before and after caps were imposed, and to further learn what general lessons, if any, this situation can teach.

The Perception of Users

From a consumer standpoint, caps and high thresholds are generally more appealing when their properties are knowable and predictable. Additionally, user behavior may be impacted substantially by incorrect understanding of contractual obligations or data use. Thus, the questions about the effects of caps can only really be answered if we understand what users know and think.

Policies for caps and thresholds should be concerned about user understanding because many household surveys find rather poor knowledge of speed/usage of own broadband and applications (see, e.g., the Pew Surveys[14]). In addition, there are changing norms for software usage and users may have a limited ability to understand the typical GB per hour of use of an application. Application and service owners bear some responsibility here as well, as they can make efforts to understand their own efficient and inefficient use of network resources, and its cost.

Lack of user understanding of how many GB may be used by applications could lead to two sub-optimal scenarios: (1) Users could underestimate the amount of data consumed and exceed their monthly data allotment, thereby incurring penalty fees or unanticipated upgrade charges; (2) Users could overestimate the amount of data consumed, thereby dissuading them from using Internet-delivered services even though they are well below their cap thresholds.

The history of unlimited dial-up can possibly explain some of the lack of user understanding of data use. The lack of limitation (i.e., unlimited use) is usually regarded as better for users than the presence of a limitation (i.e., a cap on use). Some commentators perceive an association between the lack of unlimited pricing and the lack of competitive alternative. In part, one of the most prominent historical examples reinforces the perception, namely, AOL's experience moving from usage-based pricing (specifically, metering of hours of use) to unlimited contracts. This change came about in response to competitive pressure.[15] Hence, in the minds of some commentators the increasing use of usage-based pricing with thresholds is affiliated with the decreasing use of unlimited plans, which, in turn, is presumptively affiliated with a decrease in competitive alternatives.

Unfortunately, much information about user understanding of caps and thresholds is missing. Some open questions that could be useful to answer: Do users have an ability to measure their

---

[14] See the Pew Internet and American Life Project. http://www.pewinternet.org/.

[15] The perception partly arises from the reminiscing many years later. The CEO for AOL at the time, Steve Case, states that AOL had studied the potential switch for quite some time, but not acted on it because management could anticipate a difficult transition. Competition eventually forced his hand. Said Case, "It came to a head over a weekend as Microsoft announced they were offering MSN on a flat rate basis, and it was clear they were planning to steal a lot of market share from AOL. So I decided within hours of their announcement that we had to match them, and the company worked throughout a weekend so we could make an announcement." See http://www.quora.com/AOL/How-did-AOL-make-the-decision-to-go-to-an-all-you-can-eat-pricing-strategy/. For a longer account of these events, see Swisher, Kara, *1998, aol.com: How Steve Case Beat Bill Gates, Nailed the Netheads, and Made Millions in the War for the Web*, Random House; New York.

own data use in real time? While some tools for aiding user measurement are beginning to emerge, how widely are they used and are they effective? What is the accuracy of some typical data meters?[16] Can users measure own usage by application? If so, how to encourage their use? Can users manage to monitor their use in households with multiple users and multiple devices?

The move from unlimited data to capped plans in wireless suggests some users can adjust over time to caps. However, it is difficult to predict whether that experience would carry over to wireline households, with its different applications, and in particular, whether households where multiple users of different ages occupy the residence will be able to adjust to a communal limit. However these questions of user experiences and ability to control raise questions about whether caps or thresholds that are set too low could lead to a world where the average user carefully monitors her bandwidth use, rather than leaving the average user well enough alone while only forcing "extreme" users to make changes to their use.

This topic also has implications for common notions of fairness. Typical users may be paying the same price for their Internet access as heavy users. Caps also need to be updated to match current usage patterns in order to continue to only impact "high users." From an ISP's perspective, someone who uses a steady and moderate stream of data is very different from someone who uses heavy data at peak moments of heavy use of capacity. Yet, a threshold pricing scheme hits them the same.

Another equity concern from the user perspective has to do with some models of steady data use, such as for medical purposes, which also can have implications for peak load and non-peak load use. These questions require more information about peak load pricing, a topic we take up below. For the time being, we defer more discussion.

To conclude in a similar manner to the previous section, this topic may require future monitoring, especially given the importance of consumer education to user perceptions of caps and thresholds. It is not yet apparent whether the issues in this topic are a transitory or permanent concern. The experience of ISPs with providing customers with tools to monitor or control data usage could also be valuable to insights about the perceptions of caps by consumers.

User Control

If users do not have enough control over their data usage to adequately respond, even if well informed, to caps and thresholds set by ISPs, "punishment" of users by caps or thresholds may become a problem. For example, data-intensive video commercials are increasingly being embedded in web pages by edge providers. Automated nightly/weekly updates of software are also increasingly common from software vendors. In addition, most users operate software over which the user has little control.[17]

---

[16] *See, e.g.*, Stacey Higginbotham, Feb. 7, 2013, "More bad news about broadband caps: many meters are inaccurate," *GigaOM*, http://gigaom.com/2013/02/07/more-bad-news-about-broadband-caps-many-meters-are-inaccurate/, accessed May, 17, 2013.

[17] See e.g., Peter Sevcik, 2012, "Empowering Internet users to manage broadband consumption," Netforecast, http://www.netforecast.com/wp-content/uploads/2012/06/NFR5109_Empowering_Internet_Users_to_Manage_Broadband_Consumption.pdf, accessed April 28, 2013.

Conversely, some available tools -- today used by some sophisticated users -- allow ad-blocking and other user-traffic management. Ad-blocking and flash-blocking tools are the methods most commonly discussed in online forums.[18]

User control also plays a role in discussions about overage charges. Overages only arise when a threshold is exceeded, and actual charges can depend on specific details about how overages are enforced.[19]For many users there is only downside as that threshold becomes closer. Do households consider that monitoring burdensome, particularly multi-dweller households? Do multi-dweller households perceive the monitoring as a hassle or perceive the increased uncertainty in billing as a burden? There is not enough experience yet to suggest how to characterize most households.

The working group did not further explore this topic. This issue seems largely irrelevant for the average user, as few users are affected by caps, as a practical matter. In addition, many issues in user control are too small to matter, and if they become a problem, providers typically have conversations with users, and offer amnesty. This includes issues linked to several phenomena, such as automated syncing, spam, denial of service, and compromised machines that send out messages as part of denial of service attacks. Generally speaking, the committee did not perceive these issues to be big at this time.

This may change over time. If data use grows without a commensurate increase in caps, these concerns may become urgent for policy deliberation. If this occurs, a more accurate labeling system for software applications and monitoring system that take into account caps may be a way to educate users and increase awareness of the necessity of controlling bandwidth use.

The Perception of ISPs

ISPs generally explain the use of thresholds (caps) as providing a simple pricing mechanism for matching demand for bandwidth consumption with purchasing behavior. ISPs view pricing and product choices as consumer options that are just as important to the delivery of Internet services to end users as content or technical innovations in those services.

Speed tiers also match demand for bandwidth, and most ISPs correlate speed tiers with usage thresholds. Suppliers argue that UBP with a few thresholds balances the efficiency of metered pricing without creating the stress or mental costs associated with such metering. Thus, suppliers emphasize that UBP with a few thresholds, or some forms of tiered pricing, provides a measure of bill stability, predictability, and "peace of mind" to the vast majority of consumers relative to more linear usage pricing (i.e. metered, or per KB/MB/GB, or finer-grained use tiers).

Depending on how it is structured, UBP can also enable additional lower-cost broadband plans to be offered to consumers, spurring adoption or better meeting the underserved demand from the

[18] For example, see the second comment at http://www.dslreports.com/shownews/Why-is-ATT-Capping-DSL-Users-but-Not-UVerse-Users-123692, accessed April 29, 2013.

[19] Some care is required in drawing sweeping conclusions without precise data. For example, in plans being trialed by Comcast (at the time of this writing) a user must exceed a threshold for three months in a twelve month period before overages are imposed.

low-end of the market. ISPs can afford to offer lower usage plans at a lower price point (e.g. Cable One's Economy plan): they do not add as much to the aggregate bandwidth demand for the ISP. This is one approach to manage long run bandwidth-sensitive costs.

In this sense, UBP generally serves two functions. It may affect a small number of users who use large amounts of resources. It also may shape the use of resources among the vast majority of users. As the tables showed, there are examples of ISPs pursuing policies that lend themselves to each interpretation in wireline broadband today. However, the most common so far is the use of UBP to limit a small number of users who use a large amount of resources.[20]

Beyond these generalities, more detailed analysis of the issues from a supplier's perspective falls into three categories: how to arrange prices so "high end users" pay for the additional investment they use (i.e., price discrimination in the economics literature), managing network growth (e.g., managing long run capacity investment) and managing instantaneous congestion (e.g., managing peak load pricing). The report summarizes each of these in turn.

UBP and price discrimination

Generally, in a high fixed and high sunk cost setting (such as network provision), usage based pricing is about raising revenue over incremental costs and recouping substantial fixed costs. This is generally called the economics of non-linear pricing, or price discrimination in common economic parlance.[21]

The economics literature on price discrimination provides two motives for UBP: (1) associating higher prices with higher costs and higher willingness to pay, while (2) avoiding the potential losses when some users do not buy at all. Such association can come closer to common notions of fairness and also reinforces the incentives to save on costs by showing users the price of inputs.

The Order has already made clear that usage-based pricing ensures that lighter end users are not forced to subsidize heavier end users. Charging distinct prices aligns incentives to encourage efficient use of networks. The Order also has made clear that the FCC will continue to monitor the marketplace. Thus, as the marketplace continues to develop, presumably the FCC will take these issues into account in its decision-making.

Managing Network Growth

If measurement and transparency issues were satisfactorily addressed, could a cap or threshold at a high end of downloading (e.g., less than 1% or 2% of households) reduce data use? There is little evidence (outside of Canada, as noted), so it is difficult to judge. The answer is necessarily speculative.

---

[20] We note the interesting contrast with the use of UBP in wireless contracts, where it is much more common to use UBP to shape the use of resources among the majority of users. This difference motivates open questions about why the difference arises, and what lessons can be learned from those differences.

[21] A side note about vernacular interpretation of economic terms: The word "discrimination" has a pejorative meaning in common language, though none is meant in the economics literature on price discrimination.

Generally speaking, it is thought that a data cap (in this context, a threshold with discrete changes in speed) can incentivize those near the cap to behave differently. If so, then a household that uses much more than a typical user can build more efficient usage into its own network or decrease its usage upon nearing the cap. Some of these changes may not interfere with normal Internet usage by merely optimizing bandwidth heavy behavior– for example, users streaming video footage could alter the use of uncompressed HD cameras streaming to the Internet 24x7, when on-demand will do. Users also can reduce use of peer-to-peer servers, e.g., BitTorrent, substituting partial uptime for full time. Users who run servers out of the house on a household contract could switch to a business offering that better matches needs and usage. Also, as discussed below, potentially users can take advantage of bandwidth efficiencies as they become available from edge providers. However, households that are already using bandwidth efficiently may be forced to make changes that do impact day-to-day usage.

Access providers also ask whether data caps and related means of linking price to use can encourage edge providers to innovate more efficient means of delivering their services. There is a perception that data caps and usage-based billing are not potential barriers to entry but, rather, potential drivers of greater efficiency in the delivery of edge services. They point to the incentives on Netflix and other edge service providers to innovate their services, for example, Netflix improving efficiency in Canada and licensing innovative technology like EyeIO.[22] Access providers also raise questions about the extent to which prices are misaligned and resources are misallocated because all the obligations for carriage of content is passed onto consumers (and the ISP) by edge providers. (Edge providers have a different perspective, which is discussed below.)

At most, we can draw a tentative conclusion. Over the long run a data cap or a UBP threshold can help manage network growth if users and/or edge service providers respond to the cap or threshold with less or more efficient data use; a carrier would then incur less costly operations and may be able to make less expensive infrastructure upgrades over longer periods. However, this conclusion is mostly theoretical – there is no quantitative data to suggest to what extent how much long run costs increase with growth in use or how much of a difference carrier contributions to provisioning have made to growth over time. Both the broad and specific questions cannot be answered because there is no quantitative evidence – to accept or refute – propositions about how caps and thresholds shape usage.

Managing Instantaneous Congestion

Generally speaking, instantaneous congestion management is not a stated rationale behind use of tiers, metering, or caps. There are other techniques in TCP/IP to address congestion caused by unexpected demand, outages, or major traffic shifts. Caps provide no direct incentive to heavy users to reduce traffic at peak times because there is no differential pricing across time periods. For example, monthly caps generally count traffic from the middle of the night (when traffic in general is low) against a cap.[23]

---

[22] See e.g., Janko Roettgers, 2012, "EyeIO: Netflix's secret weapon against bandwidth caps?" GigaOm, Feb 1, 2012, http://gigaom.com/2012/02/01/eyeio-video-encoding-netflix/, accessed April 28, 2013.

[23] There has been some experimentation with time-sensitive lifting of cap restrictions. See for example, this description of a satellite broadband provider's recent policy. http://www.dslreports.com/shownews/Exede-Caps-Lifted-For-Overnight-Use-120776.

However, if there is a rough correlation between total use and peak use – i.e., the largest total users over the month are also the biggest users at peak moments – then a data threshold might have some of the properties of a peak load pricing scheme by inducing a large data user to reduce their data usage. This is an open question, as there is little public analysis of the correspondence between data consumption and bandwidth usage.[24] In addition, there is little experience with alternative arrangements, as many ISPs do not perceive users calling for the option to manage data use over time.

There is no evidence, one way or another, that caps leads heavy users to reduce activity at peak moments any more than at any other moment. It would be illustrative to see if there are systematic differences between usage in the United States and Canada because of the imposition of caps and thresholds. Again, no particular data speaks to this specific question one way or another, or to the broad questions motivating it. There has not been much experience with peak load capacity management thresholds for users. Historical experience with peak load management suggests the timing for data usage and peaks would shift, but there is no evidence to suggest which applications will shift their usage patterns, or by how much they would shift them.[25]

Perception of Edge Providers[26]

A data cap or high threshold from broadband provider can shape other providers of services in broadband ecosystem, e.g., entrepreneurs who provide applications, build web pages, and operate other services in the cloud. Edge providers are concerned that a widely used cap reduced – rationally or irrationally – demand for data-intensive services and reduced entry of new data-intensive software firms, decreasing the commercialization of innovation. This concern is partially motivated by Netflix's example in Canada, which illustrates the phenomenon when a cap does bind.

Some of the power of data caps to affect edge providers that serve video or other high bandwidth media content might be offset by improvements in codecs. A codec encodes a data stream of signal for transmission, storage or encryption, and decodes it for playback and editing. (The word is a portmanteau of COder and DECoder.) There are many codecs in use today. Would improvement in codecs – i.e., to higher resolution using fewer resources with more efficiency – occur regardless of the presence/absence of caps? ISPs argue that edge providers have incentives to improve codecs when faced with caps and high thresholds. The alternative view argues that improvements arise for largely exogenous reasons, and have little relationship with the policies of ISPs.[27]

---

[24] One of the earliest studies of the correspondence between data consumption and bandwidth usage examined one ISPs traffic in 2011. It found a small correlation, not consistent with using caps to manage bandwidth. It is at http://www.fiberevolution.com/2011/11/do-data-caps-punish-the-wrong-users.html.

[25] For example, AOL experimented changes in pricing for different times of the day in order to save on phone line costs, and experienced changes in the time of day in which the "peak" usage occurred.

[26] As elsewhere in this study, we focus on the perception of "Edge Providers," as in the Order, rather than focusing on other groups of providers, such as "over the top providers," or "application service providers."

[27] The working group noted that parallel arguments take place in wireless applications.

Once again, these questions are necessarily speculative, as caps do not yet bind most households in the US, and, at present there is no decline in the demand for data-intensive services. In addition, as noted above, the experience with data limits in Canada has not been measured, so there is no data to assess the impact the caps had in that setting.

It is unclear how much entrepreneurs target already-data-intensive users. For such open questions, it is also important to recognize an asymmetry between the perspectives of edge providers and ISPs – namely, what is small to an ISP may be large to an edge provider. For example, thresholds or caps applied to a small number of households for an ISP, such as 10% to 20% of access users, can have substantial impact on the business of edge providers. A small fraction of customers to an ISP can be a large fraction of demand to a provider of data-intensive services. Fear and uncertainty could exacerbate any response, which appears to have occurred in Canada. Hence, the answer from an edge provider to these open questions could diverge from the answer from an ISP to the same open questions.

Edge providers also express a different perspective on the effects of data caps on their incentive to innovate more efficient means of delivering their services. They stress that caps could impact the deployment of new innovative services and competitors because caps disincentivize the use of more data-intensive applications. For example, in 2012, a Sony executive suggested that the company was holding off its release of an Internet video service because of ISPs data cap implementation practices.[28] Edge providers also stress that the services provided by Internet applications and websites create the value from the broadband access product offered by ISPs. Edge providers do not deliver data unless it is requested by the customers of ISPs. ISPs have an obligation for carriage of content.

We have noted elsewhere that the user response to a data cap could be exacerbated by the absence of widely used measurement tools. Here too the perspective of an edge provider may differ from that of an access provider. If users knew the "data-intensity" for various applications, they could use that information to measure the incremental contribution of each application to additional capacity use and, accordingly, adapt their own use. So there may be a consumer information dimension to this topic. For example, many edge providers offer streams of content at multiple bitrates and detect connection speed to show users a higher or lower bitrate. It's possible that edge providers could experiment with charging different prices for streams with different bitrates. What can be learned from experiments with such programs in mobile and low-bit-rate DSL?

These questions may become salient at some point for entrants who might anticipate growth in data use among US households. At what point do these concerns become urgent? If so, whose responsibility are they?

Specialized services and edge providers

In some settings, an ISP is vertically integrated into the provision of services that substitute for services a user may access over the public Internet. Thus caps may provide a method for

[28] Timothy B. Lee, May 2, 2012, "Sony: Internet video service on hold due to Comcast data cap," *Ars Technica* http://arstechnica.com/tech-policy/2012/05/sony-warns-comcast-cap-will-hamper-video-competition/, accessed May 17, 2013.

differential treatment of traffic or partners' traffic in order to favor certain applications provided by the ISP, like Voice Over IP (VoIP is a low-bandwidth application, in general).

Many aspects of this topic have been discussed by the Specialized Services working group, and we do not seek to replicate those findings here. That group has discussed questions related to incentives to build specialized services, different traffic metering to reflect different costs, difficulties with benchmarking performance in specialized services, and the different needs of distinct applications.

Here we focus on one key concern for competition policy. In general, competition policy is concerned about situations where one firm provides a service and also controls aspects affiliated with the cost, performance, and user-experience in a competing service. In public conversation this concern is often framed as a metaphor about the slope of the pitch: Does a cap or threshold tip the playing field by slanting consumers to an ISP or another online supplier? Said another way, what is a "level playing field" when a specialized service competes with an edge provider attempting to sell services that operate over the public Internet?

Despite the generality of the concerns, the answers are not sweeping or general. The specific details of this situation play an important role in determining appropriate policy. These concerns arise in a setting where managed service and Internet service use similar infrastructure, and the threshold or cap does not apply to a managed service but does apply to a range of arguably substitutable services. In such a setting, there is one set of prices and conditions for broadband service and another for the specialized service. Users pay a different price for each and have a different experience. Data caps may play a role in the prices users face and the experience they have between the two services.

This is another place where the ISP's perspective and the edge provider's perspective diverge. To see the divergence, it is useful to contrast these perspectives side-by-side.

From an ISP's perspective, since limitations do not apply to any but a small percentage of users, there is plenty of headroom for growth in competing services today and tomorrow. There is a rationale for separately provisioning between the specialized and non-specialized services, usually to achieve some engineering or market objective, such as improve the quality of service (e.g., reduce user perceptions of delay). In addition, one service often has a set of regulatory requirements associated with it, and one often does not. ISPs also note that the environment should promote innovation. For example, an ISP that is also an Incumbent Local Exchange Carrier (ILEC) transitioning to Voice over Internet Protocol (VoIP) may prioritize its VoIP traffic and exempt it from any usage threshold. In these instances, that ISP's exemption of its VoIP traffic is entirely consistent with how its traditional telephone service traffic has always been treated and should not be counted toward a cap. Any contrary conclusion would create a disincentive for the ILEC to migrate to IP and potentially stifle that migration.

From the perspective of an edge provider, similar services compete, using similar capacity, and the edge providers are providing innovative services. However, one has a threshold – say, from Hulu, Netflix, YouTube, Crackle, and competitors – and the other does not – from the ISP. The key concern is whether the rationale for distinct treatment of traffic in specialized services and

non-specialized services makes sense for the improvement of user performance, or is merely an excuse to put an edge provider competitor at disadvantage.

Does the concern arise when the thresholds are set comparatively high, as they tend to be for most ISPs today? The competition policy questions appear to be most salient in streaming of video services today, but may arise in services other than streaming. Similar issues may arise in home security systems and home video conferencing, for example. What is a level playing field in those cases?

It is difficult to forecast what users will want in a few years, and whether data caps will have any impact on those demands. It is also difficult to forecast what new applications edge providers will invent, what new specialized services ISPs will invent, and whether data caps will be relevant to their market experiences. There are both gains from flexible policy – to allow for new invention and the new situations created by invention – and gains from certainty – to allow edge providers and ISPs to plan for long-term investments. Therefore, the situation yields no easy answers in general, and, at a minimum, merits further monitoring.

Summary

This study reviewed concerns with data caps and thresholds in the context of usage-based pricing in wire-line broadband services. The report focused on providing definitions, identifying the concerns of participants, and identifying the policy issues these raised. Many open questions emerged, and full or complete answers would require considerably more discussion.

Working Group on Economic Impacts of Open Internet Frameworks

Chair:

Shane Greenstein, Professor and Kellogg Chair of Information Technology, Kellogg School of Management, Northwestern University

Members:

Brad Burnham, Founding Partner, Union Square Ventures
Neil Hunt, Chief Product Officer, Netflix
Kevin McElearney, Senior Vice President for Network Engineering, Comcast
Marc Morial, President & CEO, National Urban League
Dennis Roberson, Vice Provost & Research Professor, Illinois Institute of Technology (representing TMobile)
Charles Slocum, Assistant Executive Director, Writers Guild of America, West

Further Reading

Atkinson, Robert C., Ivy E. Schultz, Travis Korte, and Timothy Krompinger, 2011, *Broadband in America, 2nd Edition, Where It Is and Where It Is Going (According to Broadband Service Providers), An Update of the 2009 Report Originally Prepared for the Staff of the FCC's Omnibus Broadband Initiative,* CITI, Columbia University. May 2011.

Bauer, Johannes, and Steven Wildman, 2012, *The Economics of Usage Based Pricing in Broadband Markets,* Michigan State University, December 14, 2012. http://i.ncta.com/ncta_com/PDFs/Wildmanreport_web.pdf.

Bauer, Steve, David Clark, and William Lehr, 2012. *Understanding Broadband Speed Measurements,* (August 15, 2010), TPRC 2010, Available at SSRN: http://ssrm.com/abstract=1988332

Cisco, 2013, *Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012 - 17,* White Paper, Cisco Public, San Jose, CA. February 6, 2013. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html

Federal Communications Commission, 2010. *National Broadband Plan: Connecting America.* Washington, D.C. March, 2010. http://www.broadband.gov/.

Federal Communications Commission, 2013, *2013 Measuring Broadband America February Report A Report on Consumer Wireline Broadband Performance in the U.S., Washington, D.C.* February 2013. http://www.fcc.gov/measuring-broadband-america/2013/February.

Higginbotham, Stacey, 2012, "Which ISPs are Capping Your Broadband and Why?" GigaOm, October 1, 2012. http://gigaom.com/2012/10/01/data-caps-chart/

Hussain, Hibah, Kehl, Danielle, Lennett, Benjamin and, Lucey Patrick (2013), Capped Internet: No Bargain for the American Public, The New America Foundation. http://www.newamerica.net/publications/policy/capped_internet_no_bargain_for_the_american_public

Hussain, Hibah, Kehl, Danielle, Lennett, Benjamin and , Lucey Patrick (2012), Capping the Nation's Future? The New America Foundation. http://www.newamerica.net/publications/policy/capping_the_nation_s_broadband_future

Lyons, Daniel A., 2012, "*The Impact of data caps and other forms of usage based pricing for broadband access.*" Mercatus Center, George Mason University, Working Paper 12-27, October, 2012.

Odlyzko, A., Arnaud, B. S., Stallman, E., and Weinberg, M. (2012). Know Your Limits: Considering the Role of Data Caps and Usage Based Billing in Internet Access Service. Washington, DC: Public Knowledge. http://publicknowledge.org/files/UBP paper FINAL.pdf

SandVine, 2012, *Global Internet Phenomena Report,* Revision 11-06-2013, Waterloo, Ontario, Canada.

Wallsten, Scott, and James Riso, 2010, *Residential and Business Broadband Prices Part 1: An Empirical Analysis of Metering and Other Price Determinants,* Technology Policy Institute, Washington, D.C., November 2010.

Appendix 1.

	Provider	Use Threshold – GBs	Excepted Traffic	Overage Charge
1	Comcast	min 300GB (increasing by speed tier)	XFINITY Voice or Comcast Digital Voice (VoIP)	$10 / 50GB (per tier)
2	AT&T - U-Verse HSIA	250	AT&T 3G MicroCell	$10 / 50GB
	AT&T – DSL	150		
3	TWC	None	n/a	n/a
4	Verizon - FiOS / DSL	None	n/a	n/a
5	CenturyLink - 1.5 Mbps	150	Upload	None
	CenturyLink - >1.5 Mbps	250		
6	Cox - Ultimate (100 Mbps)	400	Cox Digital Voice (VoIP)	None
	Cox - Premier (25 Mbps)	250		
	Cox - Preferred (15 Mbps)	200		
	Cox - Essential (3 Mbps)	50		
	Cox - Starter (1 Mbps)	30		
7	Cablevision	None	n/a	n/a
8	Charter - Lite & Express ( )	100	None	None
	Charter - Plus & Max (30 Mbps)	250		

	Charter - Ultra100 (100 Mbps)	500		
9	Frontier	100 / 250 in selected trial markets	None	None
10	Windstream	None	n/a	n/a
11	SuddenLink (>30 Mbps)	350	We prioritize Suddenlink voice packets in order to provide quality service to our phone customers.	$10 / 50GB
	SuddenLink (10-30 Mbps)	250		
	SuddenLink (<10 Mbps)	150		
12	MediaCom - Launch (3 Mbps)	150	None	$10 / 50GB
	MediaCom - Prime (15 Mbps)	250		
	MediaCom - Prime Plus (30 Mbps)	350		
	MediaCom - Ultra/Ultra Plus (50/105 Mbps)	999		
13	Cable One – Economy	Monthly: 1GB	0000-1200 Daily	?
	Cable One - Standard (5 Mbps)	Daily: 3GB	None	None
	Cable One - Preferred (50 Mbps)	Monthly: 50 GB	0000-0800 Daily	$0.50 / 1 GB
	Cable One - Elite (50 Mbps)	Monthly: 100 GB	0000-0800 Daily	$0.50 / 1 GB
	Cable One - Premium (10 Mbps)	Daily: 5GB	None	None
	Cable One - Ultra (12 Mbps)	Daily: 10GB	None	None

14	FairPoint	None	n/a	n/a
15	Cincinnati Bell	None	n/a	n/a

15	Google Fiber	None	n/a	n/a

	Overage Treatment	**Cite**
1	Comcast does not have a cap or usage threshold but is trialing two usage based pricing plans: one with a 300 GB threshold and another with varying thresholds (the lowest being 300 GB) based on service tier.	http://corporate.comcast.com/comcast-voices/ comcast-to-replace-usage-cap-with-improved-data-usage-management-approaches
2	Notice after 1st month; notices @ 65% & 90% in following months	http://www.att.com/esupport/article.jsp?sid=KB409045#fbid=kiJ0SSZjH9l
3	n/a	http://help.twcable.com/html/twc_sub_agreement.html
4	n/a	http://www22.verizon.com/about/terms/networkmanagementguide/
5	"Customers will be given options to reduce their usage, subscribe to a higher speed residential plan, or migrate to an alternative business class high-speed Internet service."	http://www.centurylink.com/Pages/AboutUs/Legal/InternetServiceManagement/
6	If you do exceed your allowance, Cox will attempt to notify you by one or more methods: email, phone, or message on your computer before action is taken. We will then work proactively with you to resolve the problem. In many cases, customers are not even aware of their usage because they have an unsecured Wi-Fi network used by others or a computer virus. Cox can work with you to ensure that these issues are identified and corrected. In other cases, customers may choose to reduce their usage or switch to another plan that provides a higher usage allowance as Cox has assigned a different usage allowances to each of its Internet packages. In rare cases of extremely high usage Cox will suspend the user's service until they call Cox. In even rarer cases, Cox will terminate a customer's service if they do not decrease their usage after consultation with Cox.	http://ww2.cox.com/aboutus/northernvirginia/policies/speedsusage.cox
7	n/a	http://www.optimum.net/Privacy/AUP
8	Customers who exceed the "No Excessive Use of Bandwidth" section in the AUP may be notified by Charter that they have exceeded their monthly threshold and informed of Charter's Excessive Use policy. Charter Customer Care Representatives will help identify possible causes and offer suggested ways the customer can reduce bandwidth consumption. If the customer exceeds the "No Excessive Use of Bandwidth" policy and is notified three times in a six-month period, the customer's Internet service may be suspended after the delivery of the third notice.	http://myaccount.charter.com/customers/support.aspx?supportarticleid=2124
9	"In the affected markets, high bandwidth users (e.g. usage over 100Gb or 250Gb of data per month) are advised to either limit usage or convert to a high user service plan."	http://www.frontier.com/networkmanagement
10	n/a	http://www.windstream.com/Terms-and-Conditions/

11	After the first overage, the customer's Web browser will be directed to a Suddenlink notification page. The customer will be required to read that page, select how he or she wants to receive future notifications (by Web browser or email), enter the account number, and then save the information. From that point forward, future notifications on this subject will be sent each time an account reaches 80% of its monthly allowance and again when it exceeds 100%. Those notifications will be delivered through the means selected on the first overage, unless customers change their notification preference by visiting their Internet usage summary page at Suddenlink.net. Customer accounts will not be billed for exceeding their monthly allowance until the third overage. On the third and subsequent overages, the monthly allowance will be increased in installments of 50 GB at a cost of $10 per installment.	http://www.suddenlink.com/allowanceplan/
12	The data customers send and receive each month will contribute to monthly data usage. Speeds and usage allowances remain subject to change. Greater usage will result in additional charges of $10, excluding taxes and fees, for every increment of up to 50 additional Gigabytes used. For example, if usage exceeds the allowance by 51 Gigabytes, an additional charge of $20 will result.	http://mediacomcable.com/site/internet.html
13	If a user that subscribes to the Economy plan exceeds the allocated monthly bandwidth of one gigabyte, Cable One automatically will allocate a second Gigabyte to the user for a fee set forth in the subscriber agreement. If the user exceeds the bandwidth allocated by this second Gigabyte, then Cable One automatically will allocate a third Gigabyte to the user for a fee set forth in the subscriber agreement, and so on. This incremental allocation of gigabytes is valid only for the billing cycle during which it was allocated and cannot be carried forward. The total number of Gigabyte allocations and related fees charged to the user in the Economy plan is capped in the subscriber agreement.	http://www.cableone.net/Pages/internetaup.aspx
	[I]f Cable One in its sole but reasonable discretion determines that a customer has exceeded the Excessive Use threshold or is using the Service in a manner significantly uncharacteristic of a typical residential user, Cable One reserves the right to (a) adjust, suspend or terminate Service accounts at any time and without notice; or (b) require the user to upgrade his service level or pay additional fees in accordance with Cable One's then-current, applicable rates and charges for such Service; or (c) use any technology to be chosen by Cable One at its sole discretion to slow the user's service for purposes of conserving bandwidth.	
14	n/a	http://www.fairpoint.com/document/Residential_HSI_Terms_of_Service_tcm12-4842.pdf
15	n/a	http://www.cincinnatibell.com/customer_support/consumer_information/network_management/wireline.pdf
16	n/a	https://fiber.google.com/legal/network.html

Source: First fifteen observations accessed on February 6, 2013. Observation 16 accessed May 3, 2013.

AT&T/FaceTime Case Study

Mobile Broadband Working Group
Open Internet Advisory Committee
Federal Communications Commission
Released 1-17-2013

The Mobile Broadband group created a document explaining the facts behind AT&T's limited rollout of FaceTime on its mobile network, and included a number of different opinions on whether the limitations were appropriate.

The Mobile Broadband working group of the Open Internet Advisory Committee (OIAC) was formed to review the state of mobile broadband networks and assess how well Open Internet principles are working in practice. Although this report does not attempt to engage in any legal interpretations of the Open Internet Order, we do note that the Order [http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf] treats these mobile broadband networks differently from traditional fixed networks. While both fixed and mobile broadband providers must disclose their management practices, mobile broadband providers have greater latitude for blocking devices and applications (as long as they do not compete with the provider's own voice or video telephony services) and discriminating in how they serve traffic, in accordance with reasonable network-management practices.

The working group is investigating the tension between the goals of a free and open Internet, and the very real challenges that arise in managing mobile broadband networks. Such an investigation can easily devolve into vague discussions of high-level concepts or principles that may not be realizable in practice. To ground the discussion, the group started by considering several concrete case studies to help identify important trade-offs, principles, and other issues warranting further study, rather than trying to reach consensus on specific policy recommendations. The group explored one timely case study concerning how AT&T restricted the use of Apple's FaceTime application over its cellular data network to customers subscribed to a particular pricing plan. Video communication is widely viewed as the logical next step beyond the delivery of voice, text, and images over cellular data networks. Yet, these applications consume significant bandwidth and often have strict performance requirements, making them especially challenging for carriers to support efficiently. In the rest of this report, we discuss the specifics of the case study, analyze the high-level issues it raises, and present several possible conclusions from the unique perspectives of application developers, carriers, and equipment vendors.

AT&T and FaceTime

FaceTime is a high-quality video-calling service created by Apple for use on the iPhone, iPad, and Mac. On the iPhone, rather than operating as a separate application, FaceTime is automatically integrated into the normal calling features of the user device. A user can upgrade a conventional phone call to include video simply by pressing a FaceTime button. Originally, Apple made FaceTime available only over wireless (WiFi) connections to the Internet, and the FaceTime calling features could not be used when devices were connected to a cellular network; however, that restriction was recently lifted, in part.

In June 2012, Apple announced that FaceTime would be available over cellular data networks, though Apple acknowledged that carrier restrictions may apply. In August 2012, AT&T announced that, in the wake of Apple's lifting of its restriction on FaceTime use, AT&T would limit the use of FaceTime over its cellular data network to customers of its MobileShare plans, in which multiple devices share a single limit for total data usage. Customers with "unlimited" data plans would not be able to use FaceTime on AT&T's cellular data network. The requirement for a specific plan would be enforced directly by the device, based on carrier settings [http://support.apple.com/kb/HT1970] (such as the current data plan or other eligibility information) learned from the carrier when the device authenticates with the cellular network.

Other providers, such as Sprint and Verizon, announced that FaceTime would operate over their cellular data networks for users of all billing plans [http://9to5mac.com/2012/07/18/sprint-says-it-will-not-charge-for-facetime-over-cellular-verizon-calls-talk-premature/, http://arstechnica.com/apple/2012/09/verizon-will-enable-iphones-facetime-on-all-data-plans-unlike-att/].

Some advocates and press denounced AT&T's decision, claiming that AT&T was violating the FCC's Open Internet Order [http://www.savetheinternet.com/press-release/99480/att-blocking-iphones-facetime-app-would-harm-consumers-and-break-net-neutrality, http://publicknowledge.org/att-facetime]. They argued that AT&T was blocking an application competing with its own voice or video telephony services, and that reasonable network management practices do not include favoring one pricing plan over another.

Responding to these claims, a blog post by AT&T [http://attpublicpolicy.com/fcc/enabling-facetime-over-our-mobile-broadband-network/] argued that AT&T's policy was fully transparent, and that AT&T does not have a competitive video calling application. AT&T also argued that the FCC's Open Internet Order does not regulate the handling of pre-loaded applications (i.e., applications integrated into the device's operating system, rather than installed manually by a user). AT&T also noted that all customers can continue running FaceTime over WiFi connections to the Internet.

In September 2012, several public interest groups announced their intent to file a formal complaint with the FCC [http://arstechnica.com/tech-policy/2012/09/att-faces-formal-fcc-complaint-for-blocking-cellular-facetime-use/], arguing that AT&T's restrictions of FaceTime usage violate the Open Internet Order. In October 2012, an AT&T customer in San Francisco filed a consumer complaint with the FCC concerning AT&T's blocking of FaceTime on his "unlimited" data plan [http://www.businessinsider.com/consumer-fcc-complaint-att-facetime-2012-10].

On November 8, 2012, AT&T announced [http://attpublicpolicy.com/consumers-2/a-few-thoughts-on-facetime/] plans to support FaceTime on all of its tiered data plans for users with an LTE device, over the next 8-10 weeks. AT&T customers with non-LTE devices or unlimited data plans would still not have access to FaceTime over the cellular network. AT&T also began rolling out new billing plans to enable deaf and hard-of-hearing customers to use FaceTime.

Main Issues

AT&T's restrictions on the FaceTime application raise several interesting issues:

Pre-loaded application: Unlike many applications, FaceTime comes pre-loaded on a very popular phone. The application is immediately available to all users of the phone without requiring purchase or download, and is accessed via the core calling functions of the device. Every time a customer makes a phone call, the option of using FaceTime is immediately available. This makes it much more likely that the application would enjoy large-scale adoption very quickly. In addition, simultaneous use of the application (say, by spectators at a sporting event) could overwhelm the available radio network capacity, with its finite spectrum. In contrast, applications that require a manual download typically see lower penetration, even for popular applications that can be downloaded free of charge. For example, while around 75 million iPhones were sold in 2010, Skype was downloaded to only 7 million iPhones, resulting in less than 10% penetration [http://www.statisticbrain.com/skype-statistics/]. The rapid availability of FaceTime is said to be a particular challenge for AT&T, which historically has a much larger penetration of Apple iPhones among its customers, compared to other carriers [http://news.cnet.com/8301-13579_3-57492508-37/iphone-owned-63-percent-of-smartphone-marketshare-at-at-t/]; today, more than half of AT&T's cellular data-network subscribers use an iPhone.

High bandwidth requirements: Cellular data networks have limited capacity, particularly in the "upstream" direction from user devices to the Internet; as such, carriers must carefully manage the shared "up-link" bandwidth to ensure reasonable performance for all users. While most content-delivery applications primarily impose load on the "down link," high-quality, video-telephony applications (like FaceTime) typically generate a large amount of traffic in both directions to deliver high-quality video to both participants in a video phone call. The quality of a multimedia application depends on the available bandwidth. Most popular applications adapt automatically in the presence of congestion, to decrease the quality of the audio or video stream to share bandwidth fairly with other applications. For example, data from Skype suggests that 128-300kbps is required for a standard video call [https://support.skype.com/en/faq/FA1417/how-much-bandwidth-does-skype-need], whereas various online reports suggest that FaceTime consumes around 100kbps - 1000kbps [http://www.tested.com/news/254277-why-is-att-doing-you-a-favor-by-blocking-facetime/, http://www.padgadget.com/2012/06/20/concerns-about-facetime-over-cellular-will-you-max-out-your-data- limits, http://appadvice.com/appnn/2012/10/its-pretty-stupid-ridiculous-how-much-data-netflix-uses-over-lte, http://www.nokiasiemensnetworks.com/system/files/document/smart__labs_-_facetime_over_cellular_in_iphone_ios6_final_0.pdf], consistent a limited set of measurements conducted at Bell Labs at the request of this working group. It therefore seems to be the case that FaceTime currently consumes on average 2-4 times more bandwidth than a similar Skype video call. It is important to note that there is no fundamental reason why FaceTime could not adapt to congestion the same way as other applications, and the way FaceTime behaves in the presence of congestion may easily change in the future.

Staged deployment of new applications: Rapid adoption of a new application might lead to large and unpredictable changes in the traffic load on a cellular data network. Carriers may want to

start with a limited trial deployment of a new application to better understand its effects before wide-scale deployment. This can provide measurement data and operational experience that carriers and application developers can use to make the most effective use of limited resources, or to identify appropriate policies for sharing resources with other applications. The AT&T/FaceTime case study raises an interesting question of whether or not restricting usage to customers of a particular pricing plan is a good way to limit (i) the number of users in an initial deployment (i.e., to users of a particular plan) or (ii) the total volume of traffic (i.e., by denying access to users with unlimited data plans), and what other alternatives might exist.

Application management on the device vs. the network: A carrier can block an application by discarding the packets it sends or receives; alternatively, a device such as a smart phone can prevent users from running a particular application, thereby keeping the traffic from ever reaching the network. In the AT&T/FaceTime case study, the usage of FaceTime on AT&T's network was limited directly on the device, rather than inside the network. An interesting policy question is whether it matters where an application-management decision is enforced, and which organization decides what policies to place on an application's use. In some cases, the creator of an application may want its users to enjoy unfettered access to the application, but in others the application developer may prefer to limit usage to ensure that supported users enjoy good performance; distinguishing between these two situations is surprisingly difficult. In this case, Apple and AT&T have not commented on which organization initiated the restrictions, and whether or not this was a collaborative decision.

These issues demonstrate the subtle trade-offs that arise in determining whether restricting FaceTime usage over AT&T's network constitutes blocking and/or reasonable network management.

Summary Opinions

Different members of the working group came to different opinions about the restriction of FaceTime usage on AT&T's network. Generally, the working-group members agreed that blocking applications runs the risk of discouraging innovation, but that carriers also need effective ways to manage the limited resources in cellular networks. This led to three main opinions about AT&T's decision to restrict customer access to the FaceTime application over its cellular network, presented from the perspectives of different parts of the mobile broadband ecosystem -- application developers, carriers, and network equipment vendors. These opinions convey the conclusions of advocates for these perspectives among the working-group members, but do not attempt to fully represent each community.

- From the perspective of application developers:

AT&T did not choose the optimal approach by blocking access to the FaceTime application for customers on certain data plans. By singling out one popular application, the door is opened for carriers to block lawful use of applications, require customers to upgrade to potentially costlier, limited plans, and justify their actions by claiming to be engaged in reasonable network-management practices. Unfortunately, blocking a specific application for a large number of users on certain pricing plans, instead of managing the congestion that application and others might cause, sets a precedent that could have very negative consequences for the vibrant market for mobile applications. Allowing application blocking means that no developer could be sure that

his or her mobile application will be able to reach customers. If a carrier can block an application entirely at its discretion, investors will have to consider a new risk in addition to the normal risks faced by any start up. Unlike technical risk, financial risk, or organizational risk, the risk of being blocked cannot be mitigated. The existence of that risk will limit the investment available to applications developers, limiting the number of applications created, slowing innovation, and limiting consumer choice.

AT&T may have chosen to block FaceTime because it was a simple way to manage the potential congestion that could have occurred if the application were widely used. The carrier may have chosen to block FaceTime because it was concerned that broad use of a high-bandwidth data application by users of unlimited pricing plans would impact its profitability. Managing congestion and profitability are legitimate objectives for AT&T, but furthering those objectives by blocking specific applications is not the way to do it. There are many ways AT&T could have managed the roll out of FaceTime over cellular without taking the kind of application-specific action that harms applications developers and ultimately consumers. For example, AT&T could have instituted rate-limiting of individual customers, applied in a neutral manner, to limit congestion. Rate limits could be imposed at peak times or in response to congestion. In the medium- or long-term, AT&T could more aggressively scale up network capacity or apply other bandwidth-management techniques (such as WiFi offload) in localized hot spots where FaceTime and other high-bandwidth applications create congestion problems. AT&T can also charge users for the amount of data they consume, independent of the application. We recognize that these approaches require AT&T to deploy the technology in the network to actually manage the network, or to make the investment to market a new pricing plan to consumers. We understand that blocking FaceTime may be simpler and cheaper than deploying new network-management technology, increasing capacity, or changing pricing, but blocking a specific application chills investment, harms application developers, and reduces consumer choice. That is too high a price to pay when other alternatives are readily available.

In short, network management should focus on the underlying conditions that cause degraded performance of the network and address those conditions with solutions that optimize performance in a neutral manner for all users and applications. Such approaches -- indeed, all aspects of traffic management and engineering -- may require advanced planning to ensure that they are available when network conditions require them, but that fact makes them no less appropriate from a technical perspective. Application-agnostic network-management approaches should be considered and exhausted before application-specific approaches are even considered on a temporary basis, and customers should be able to have their choice of applications without having to change their data plans. Giving customers choice includes the option for user-controlled quality of service, where users decide to favor traffic from one application over another, in allocating whatever share of network bandwidth they receive from the carrier.

- From the perspective of carriers:
Given the bandwidth-intensive nature of the FaceTime application and AT&T's significant base of iPhone subscribers, AT&T has good reasons to be concerned about the potential for FaceTime to cause a focused, or localized, overload condition in its network. AT&T's approach of enabling FaceTime on Wi-Fi and on cellular for shared data plan subscribers is a reasonable way of managing the risk of network congestion. As data about FaceTime usage becomes available and

as its network evolves, AT&T has indicated that it may further expand the availability of the application. In fact, AT&T has already expanded the availability of the application to users with LTE devices on tiered service plans and on new custom plans for the hearing impaired.

AT&T's approach reduces the probability of a focused overload of its network due to FaceTime usage. By requiring a usage-based plan to access FaceTime over the cellular network, AT&T's approach both encourages use of the FaceTime service in a manner that is less likely to adversely impact the experience of other users on the network, and manages the number of subscribers that are likely to use such a bandwidth-intensive application. Usage-based data plans provide an incentive for users to manage their consumption of network bandwidth, and ensure that heavier bandwidth users pay a proportionate amount for their usage when compared to lighter bandwidth users. Unlimited data plans provide no incentive to users to manage the data consumed by bandwidth-intensive applications. Unlike some of its competitors, AT&T continues to offer unlimited data plans to existing subscribers to those plans, even when those subscribers upgrade to a new subsidized device. Since some carriers mandate that subscribers switch to a shared data plan when upgrading to a new device, AT&T's approach gives customers more flexibility than some of its competitors in choosing pricing plans and services that meet their needs. AT&T's announced expansion of FaceTime availability to LTE devices on individual tiered plans recognizes the increased capacity of its LTE network which, unlike its UMTS network, is not currently carrying voice calls, thus balancing the overall service quality for all of its customers.

While critics of AT&T's approach have described possible alternative approaches to the situation, none of the alternatives would effectively address AT&T's concerns. AT&T is aggressively expanding its cellular network capacity, and its devices are configured to support offload of data traffic to Wi-Fi networks where possible. AT&T currently operates over 30,000 Wi-Fi hotspots freely available to its data plan subscribers. While some have proposed rate limiting subscribers during periods of congestion, this approach is problematic for two reasons. One reason is that dynamic rate limiting is a complex mechanism that is not currently supported by wireless standards and vendor equipment. While dynamic rate limiting might be an option in the future, it is not an option that is available to AT&T today. The second reason is that dynamic rate limiting has the potential to degrade performance for both FaceTime and other applications. As a result, rate limiting may lead to more user dissatisfaction than AT&T's approach. This does not rule out dynamic rate limiting as a potential solution. However, it illustrates the complexity of providing good quality mobile broadband services.

While some have argued that AT&T's approach may adversely affect innovation, this risk can be mitigated by application developers by working cooperatively with carriers to build applications that do not risk harm to the network. In the case of FaceTime, the company developing the application built a mechanism into its operating system that enables operators to require certain plans. Other non-US carriers have used the same mechanism. Apple's page at http://support.apple.com/kb/ht1937 shows the carrier-by-carrier breakdown of features supported by carriers world-wide. This specific example does not support the "chill to investment" argument, as the dominant player allowed its offering to be managed, which is rather different from a new entrant struggling to break in to a market.

In making these types of decisions, carriers are weighing multiple factors and taking competitive risks that may or may not succeed in the marketplace, but the marketplace can and should determine the success of these approaches. These decisions and the set of available techniques are not static and cannot be proscribed or regulated with any reasonable degree of applicability or validity over time.

- From the perspective of network equipment vendors:
Applications supporting real-time, two-way video calling such as Skype have become increasingly popular (more than 100 million logins/month and 30 million simultaneously active calls [http://www.statisticbrain.com/skype-statistics/]) and this popularity has increased with the availability of mobile clients for these applications. Given the significant additional bandwidth requirements of video sessions over voice calling, encoding the video frames at lower bit rates and the ability to adapt to changing network conditions such as the available bandwidth is key to the successful deployment or use of such applications. This is particularly true for mobile networks which represent a highly constrained and shared resource in both the uplink and downlink directions. For these reasons Skype utilizes adaptive session control techniques to constantly adjust the bit rate of the video stream transmitted between the two endpoints.

Apple's Facetime application is targeted to the same video calling market segment, but as noted above does not seem to adapt as readily/aggressively to changing network conditions. To illustrate the additional potential consumption compared to Skype usage, consider the following: if, as stated above, 10% of iPhone users were Skype users. When one compares this to the 100% of iPhone users who have access to the Facetime client and the at least 2x additional bandwidth consumption by the iPhone Facetime client compared to the Skype client, it is reasonable to conclude that the total network bandwidth usage (across all users and sessions) of Facetime could be as much as 20x higher than that of Skype, for operators who have a significant proportion of iPhones in their network.

In this context, it is reasonable to conclude that AT&T, with the largest number of iPhone users and largest fraction of their subscribers as iPhone users would have particular concerns about the load that the Facetime application would put on their network, with the potential to significantly degrade the available bandwidth for all other applications. Moreover, the concern would be most prevalent with respect to the most scarce resource -- the cellular network (which typically has ~20Mhz of spectrum compared to the more than 100Mhz of WiFi spectrum [http://en.wikipedia.org/wiki/List_of_WLAN_channels]) -- which is also the resource for which users have the highest service expectation. Given this, there would be a clear imperative to manage the usage of FaceTime application on AT&T's cellular network, with the option for unmanaged usage of FaceTime over their network. This is precisely the behavior that AT&T exhibited by limiting the usage of FaceTime to only a subset of their pricing plans, whilst making FaceTime available to all users over the WiFi interface. As such, it is reasonable to conclude that AT&T was trying to employ reasonable network management to the use of FaceTime over their network, albeit it in a relatively crude form.

It is interesting to contemplate whether there are alternative means by which the usage of FaceTime could have been managed in a way that would have made it available to all cellular users but in a scalable way. Clearly, if FaceTime was similar to Skype in terms of its bandwidth

utilization per session, or having the same device penetration (available on 10% of iPhones), no specific network management would have been required relative to that applied for Skype. Therefore an alternative approach would have been for AT&T to work with Apple to improve the bandwidth adaptation capabilities of the FaceTime application. Another alternative approach could have been to rate limit the usage of FaceTime in the network both on an individual session basis (per user), and an aggregate basis (total bandwidth allowed for all FaceTime users) using rate limiting techniques employed by some operators when usage caps have been reached, or for service plans that explicitly exclude usage of certain applications. Last, a non-application-specific rate-limiting approach could have been employed whereby the peak bandwidth usage by each user was limited when the network was congested. These approaches would have been reasonable and preferable in terms of the universal applicability and equanimity of the solution. It is important to note, however, that these alternative approaches may actually have resulted in a less satisfactory experience for all FaceTime users, or across all applications being used (for the non-application-specific approach), in contrast to the approach that AT&T took which likely resulted in a more satisfactory FaceTime experience, but for a subset of users. In other words, non-application specific approaches can appear 'fair' as they apply a 'one size fits all' philosophy whereby all users receive the same treatment for all applications. But, in some cases, and at some points in time, users may have a preference for a certain application (e.g. a FaceTime session for an important call) and would prefer it to be prioritized over other internet-based services when the network is congested.

Conclusion

The three summary opinions capture different perspectives, with some overlapping points and differences in emphasis. Most members of the working group agreed with aspects of all three opinions, with some aligning more strongly with one view over the others. The case study also highlights the need for future cellular networking equipment and management systems to offer greater flexibility in managing the fine-grain sharing of limited network resources. This would make it easier for carriers to limit the impact new applications have on the performance experienced by other users using application-neutral techniques.

FCC Open Internet Advisory Committee Mobile Broadband Working Group

- Chair: Jennifer Rexford, Professor of Computer Science, Princeton University
- Harvey Anderson, Vice President of Business Affairs & General Counsel, Mozilla
- Brad Burnham, Founding Partner, Union Square Ventures
- Alissa Cooper, Chief Computer Scientist, Center for Democracy & Technology
- Charles Kalmanek, Vice President of Research, AT&T
- Matthew Larsen, CEO, Vistabeam
- Dennis Roberson, Vice Provost & Research Professor, Illinois Institute of Technology (representing TMobile)
- Chip Sharp, Director, Technology Policy and Internet Governance, Cisco Systems
- Marcus Weldon, Chief Technology Officer, Alcatel-Lucent
- Jonathan Zittrain, Professor of Law and Professor of Computer Science, Harvard; OIAC Chair (ex officio)
- David Clark, Senior Research Scientist, MIT Computer Science and Artificial Intelligence Research Laboratory; OIAC Vice Chair (ex officio)

Openness in the Mobile Broadband Ecosystem

Mobile Broadband Working Group
Open Internet Advisory Committee
Federal Communications Commission

The Mobile Broadband group also created an analysis of the mobile broadband ecosystem, identifying key players and articulating their relationships.

The FCC's Open Internet Order[29] characterizes "openness" as "the absence of any gatekeeper blocking lawful uses of the network or picking winners and losers online" and indicates that the openness of the Internet promotes a self-reinforcing "cycle of investment and innovation" (p. 3). In the mobile broadband ecosystem, a variety of players have significant roles in shaping the opportunities that the Internet provides, including mobile broadband providers (e.g., Verizon, AT&T, Sprint, and T-Mobile), device vendors (e.g., Apple, Samsung, and LG), operating system developers (e.g., Apple iOS and Google Android), network equipment vendors (e.g., Ericsson, Alcatel-Lucent, and Nokia-Siemens), and application developers and content providers.

This report examines the relationships between these parties and highlights the different kinds of influence they can have over openness, broadly defined. While many of these parties are not subject to the Open Internet Order, understanding the impact they can have on openness provides a more complete picture of the mobile broadband ecosystem. Because of our specific focus on mobile broadband, our analysis inherently reflects business and technical dynamics that may differ from those for fixed broadband networks. Also, while mobile broadband networks carry a variety of traffic (e.g., downloading e-books to Kindle devices, machine-to-machine communication, connected cars, etc.), this report focuses on the general, universal service that connects end-user mobile devices to the Internet.

1. Mobile Broadband Ecosystem

The mobile broadband ecosystem is built on a seemingly "virtuous cycle," where *networks* that are fast, reliable, and widely available encourage the creation of *mobile devices* that connect to these networks, which spurs innovation in compelling *applications and content*, which in turn motivate more *users* to adopt the technology, spurring further investment in the underlying networks.

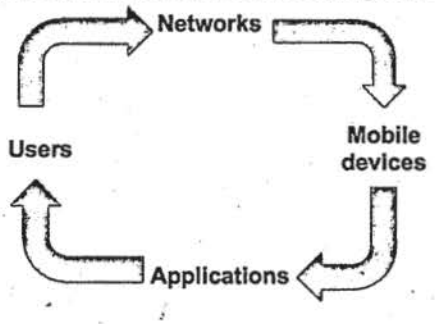

[29] FCC Open Internet Report and Order, December 2012.
http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf

Yet, the players in the mobile broadband ecosystem have complex relationships that can cause tensions that can dampen the incentives for innovation and investment. The main parties include the network (i.e., *mobile broadband providers* and network *equipment vendors*), the devices (i.e., *device manufacturers* and *operating-system developers*), the applications (i.e., *application developers*), and the *component manufacturers* who make the components used in mobile devices and network equipment.

1.1 Major Mobile Broadband Companies in the U.S. Market

In most sectors of the mobile broadband ecosystem, a small number of companies drive the market, as shown in the following table:

	Ecosystem Players in the U.S. (1Q 2013)
Smartphone vendor shipments[30]	Apple (38.3%), Samsung (28.8%), LG (9.9%), and many smaller players (< 5% each)
Smartphone OS market share (through 1Q13)[31]	Google Android (56.0%), Apple iOS (38.3%), and other smaller players (< 4%)
Mobile broadband provider market share[32]	Verizon Wireless (34%), AT&T Mobility (30.9%), Sprint (16%), T-Mobile USA (12.2%), and other smaller players (< 3%)
Radio access network equipment vendors[33]	Ericsson (50%), Alcatel-Lucent (36%), Nokia-Siemens (10%), Huawei (3%)
Application developers[34]	Many, diverse, most make < $500/month

A few main vendors lead the sectors for creating smart phones (e.g., Apple, Samsung, and LG) and the operating systems that run on them (e.g., Google Android and Apple iOS), along with some smaller players. The U.S. has four main mobile broadband providers (Verizon, AT&T, Sprint, and T-Mobile). Mobile broadband providers can acquire equipment for cellular access networks from three main vendors (Ericsson, Alcatel-Lucent, and Nokia-Siemens), with Samsung a new entrant into the U.S. LTE equipment market, and Huawei a smaller player in some U.S. regional markets. In addition, a small number of companies create most of the components used in handsets (e.g., Qualcomm and Samsung) and the components used in network equipment (e.g., Texas Instruments, Broadcom, and Freescale). In contrast, the applications sector is extremely large and diverse, with many thousands of developers

[30] Strategy Analytics, "North America Smartphone Vendor & OS Market Share by Country: Q1 2013," May 2013
[31] Ibid.
[32] Strategy Analytics, "Wireless Operator Performance Benchmarking Q4 2012," April 2013
[33] Alcatel-Lucent internal analysis of Dell'Oro data, average over the last four quarters.
[34] Source: Vision Mobile

creating applications that compete for users' attention. The app market generated more than 13.4 billion downloads and $2.2 billion of revenue[35] in the first quarter of 2013 alone. While most application developers operate at a very small scale (e.g., making less than $500 per month), half of all app revenue comes from just 25 developers[36] --- mostly major game developers such as Zynga, Electronic Arts, Rovio, and Disney.

While mobile broadband providers are typically regional or national companies, the rest of the mobile broadband ecosystem is an international marketplace. While most of the leading companies in the U.S. have significant market share internationally, some companies play a much larger role in the rest of the world; for instance, Huawei has a much larger market share in the network equipment market internationally. Historically, the U.S. was the leader in cellular deployments, but lost the lead to Europe in 2G (GSM) and to Asia in 3G (WCDMA), before regaining the lead again with 4G (LTE). The U.S. also leads the recent innovations in smart phones, mobile operating systems, and applications. Still, the manufacturing of components and handsets mainly takes place outside the U.S., and the mobile broadband ecosystem relies heavily on international agreement for technology standards. In addition, many new mobile-broadband business trends, such as the decreasing role of carrier subsidies for mobile handsets, started outside the U.S., providing a unique opportunity to analyze the effects of emerging trends.

Some companies play a significant role in multiple parts of the mobile broadband ecosystem, giving them extra influence. While industry forces often work against having a primary "vertical player" (e.g., Motorola, in earlier days), several companies increasingly play multiple roles in the mobile broadband sector. For example, the top handset manufacturer (Samsung) also sells LTE equipment, as well as the low-level components used in other handsets (such as the Apple iPhone)[37]. Huawei also sells both mobile devices and network equipment. As such, Samsung and Huawei can have a unique relationship with carriers, by having bundled offerings of handsets and network equipment. Apple and Google also have significant influence in multiple parts of the ecosystem. Apple creates devices (e.g., iPhones and iPads) that are tied to its own operating system (iOS), and also develops mobile applications that come bundled with the device. Google has the lead mobile operating system (Google Android), and also creates popular applications and, recently, mobile handsets. In the next subsection, we discuss the interaction between these and other companies in the mobile broadband ecosystem.

1.2 Complex Inter-Relationships in the Mobile Broadband Ecosystem

Each of the players in the mobile broadband ecosystem is affected by the policies and practices of the others, including:

Users: End-users identify strongly with their mobile devices, from the early Razr flip phone to the Apple iPhone. With the emergence of smart phones, users increasingly associate their entire mobile broadband experience with their device, and particularly with the operating system (e.g., Apple iOS and Google

[35] http://news.cnet.com/8301-1035_3-57578563-94/app-market-soars-with-13.4-billion-downloads-in-q1-2013/

[36] http://www.canalys.com/newsroom/top-25-us-developers-account-half-app-revenue

[37] http://www.economist.com/blogs/dailychart/2011/08/apple-and-samsungs-symbiotic-relationship